Frequency Domain Beamformer for a 3-D Sediment Volume Imaging Synthetic Aperture Sonar

Jonathan R. Pearson Magoon,[a] Matthew A. Nelson,[a] Daniel D. Sternlicht,[b]

[a] Applied Signal Technology, Inc.
20101 Hamilton Ave, Suite 150
Torrance CA 90502 USA

[b] Naval Surface Warfare Center, Panama City
110 Vernon Ave
Panama City FL 32407 USA

Abstract-A frequency domain beamforming approach is described for 3-D sediment volume imaging synthetic aperture sonars (SAS). The beamformer, designed for systems with receiver arrays oriented transverse to the vehicle, performs standard delay and sum processing for each ping in the across-track direction followed by multi-ping synthetic aperture processing of elevation planes in the along-track direction. The beamformer can be run repeatedly with varying along-track squint angles revealing aspect-dependent target features. The collection of elevation planes, independently computed by the frequency domain beamformer, can be processed simultaneously on multiple threads of execution, in order to take advantage of multi-CPU machines.

I. INTRODUCTION

This paper describes a frequency domain beamforming approach for 3-D sediment volume imaging synthetic aperture sonars (SAS). The beamformer is designed for systems with receiver arrays oriented transverse to the vehicle such as Florida Atlantic University's Buried Object Scanning Sonar (BOSS) [1] and the sub-bottom imaging SAS co-developed by Alliant Techsystems Inc. and the University of Hawaii [2].

The technique uses near-field, delay and sum processing on each ping to beamform in elevation the sonar returns onto particular angles and ranges from the vehicle. After this is complete for a set of pings, the elevation beam planes are independently SAS processed using a frequency-domain beamformer that implements the Range Migration Algorithm (RMA). The SAS-processed elevation planes are transformed onto a geo-referenced three-dimensional Cartesian grid. The RMA beamformer can efficiently be run repeatedly with varying along-track squint angles, revealing aspect-dependent target features.

Motivation for developing a frequency domain beamformer lies in its superior computational efficiency, vice the more commonly implemented time domain counterpart, and in the relative speed at which images from multiple aspects may be produced.

Section II of this paper describes the 3-D beamforming procedure, Section III describes data and results, and Section IV reviews advantages and next steps.

II. 3-D BEAMFORMING PROCEDURE

The beamforming procedure described here is demonstrated with the BOSS – a broadband sonar capable of using energy backscattered from the sediment volume to generate decimeter resolution, multi-aspect 3-D imagery of buried objects. The system employs a 3–20 kHz omni-directional projector and a 2-meter long, 40-element horizontal receiver array oriented transverse to the axis of the vehicle and embedded into two laterally mounted wings (Fig. 1). The required vehicle navigation solutions are supplied by a DVL-aided Inertial Navigation System (INS).
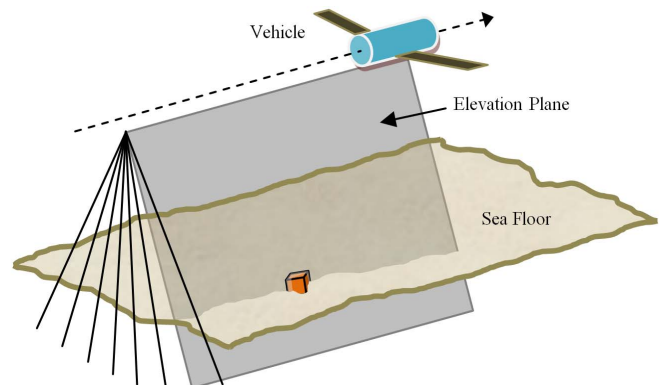

Fig. 1. BOSS SAS Beamforming Geometry. Sonar returns from each ping are near-field time-domain beamformed onto elevation planes each with an origin along the vehicle track line. The returns from each elevation plane for a collection of pings (a frame of data) are independently compensated for platform motion and are then SAS processed using the Range Migration Algorithm. The SAS processed returns from all of the elevation planes are transformed onto a Cartesian grid for data analysis and display.

Preliminary Processing

A replica of the LFM transmit pulse is used to de-chirp the raw sonar data, during which a Hanning frequency window is applied to select a sub-band of frequencies. Lower frequencies are included when bottom penetration is to be emphasized.

Report Documentation Page

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
JUN 2010	**N/A**	**-**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Frequency Domain Beamformer for a 3-D Sediment Volume Imaging Synthetic Aperture Sonar	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Applied Signal Technology, Inc. 20101 Hamilton Ave, Suite 150 Torrance CA 90502 USA	

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM202806. Proceedings of the Oceans 2009 MTS/IEEE Conference held in Biloxi, Mississippi on 26-29 October 2009. U.S. Government or Federal Purpose Rights License., The original document contains color images.

14. ABSTRACT
A frequency domain beamforming approach is described for 3-D sediment volume imaging synthetic aperture sonars (SAS). The beamformer, designed for systems with receiver arrays oriented transverse to the vehicle, performs standard delay and sum processing for each ping in the acrosstrack direction followed by multi-ping synthetic aperture processing of elevation planes in the along-track direction. The beamformer can be run repeatedly with varying along-track squint angles revealing aspect-dependent target features. The collection of elevation planes, independently computed by the frequency domain beamformer, can be processed simultaneously on multiple threads of execution, in order to take advantage of multi-CPU machines.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE	**SAR**	**4**	
unclassified	**unclassified**	**unclassified**			

Elevation Plane

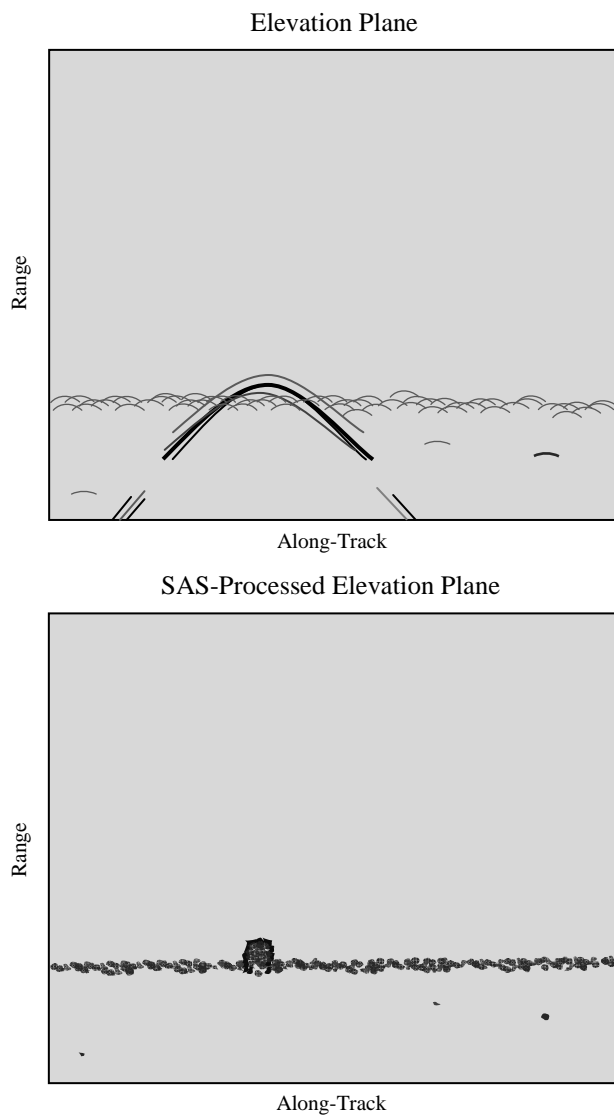
SAS-Processed Elevation Plane

Fig. 2. SAS Beamforming. Top figure shows a diagrammatic representation of a single elevation plane prior to SAS processing. The elevation plane shown here corresponds to the highlighted plane in Fig. 1. Notice hyperbolic-shaped returns for a given scatterer in the scene. The bottom figure shows the elevation plane after SAS processing. The SAS processing has focused the elevation plane in the along-track direction.

Selection of higher frequencies and wide bandwidths combine to enhance fine scale resolution in the across-track direction.

The strong specular return from the water-sediment interface

may reduce image contrast for objects in the overall field of view. Thus, a DVL-extracted measure of altitude is used to implement a range-gate to null these high intensity voxels.

Elevation Beamforming

Near-field, delay and sum processing on each ping is used to beamform in elevation the sonar returns onto particular angles and ranges from the vehicle. Sensor roll, pitch, heave (upward motion), and sway (cross-track motion) are compensated for by adjusting the range used in the coherent sum for each sonar channel. The motions that typically have the greatest effect on elevation beamforming are roll and heave since they most significantly affect the path length between a hydrophone and a location in space below the vehicle. To obtain estimates of heave and sway, the INS velocities are integrated to form position estimates. A linear least-squares fit is computed for each of the three coordinate directions. Once this best-fit line is computed, the INS-estimated vehicle trajectory is used to estimate deviations from the best-fit line at each point in time. Linear interpolation is used whenever an INS measurement or motion estimate does not fall directly on a ping time.

The collection of elevation beams is chosen to adequately sample the angular resolution of the lateral wing array. The origin of the elevation beams is chosen to be the virtual centerline of the vehicle motion (the linear fit to the computed vehicle trajectory).

Along-track Beamforming

Once a beamformed elevation plane is generated, it is ready for further motion compensation and SAS processing. To compensate for the uneven surge of the vehicle, vehicle motion extracted from the INS is used to resample the acoustic data so that all phase histories are referenced to a uniform grid (a requirement for frequency-domain SAS processing).

The elevation beams are SAS processed individually using the Range Migration Algorithm (RMA) [3]. Standard RMA routines from Applied Signal Technology's PRO*SAS*™ beamformer are used. See Fig. 2 for an illustration of the results of SAS-processing an elevation plane.

Multi-aspect Processing

In processing each synthetic aperture elevation plane, the corresponding phase histories can be delayed such that energy scattering from a particular direction is favored. This process, called squinting, is carried out by selective windowing of the Doppler shifts included in the frequency domain during RMA, effectively removing energy from unwanted look directions.
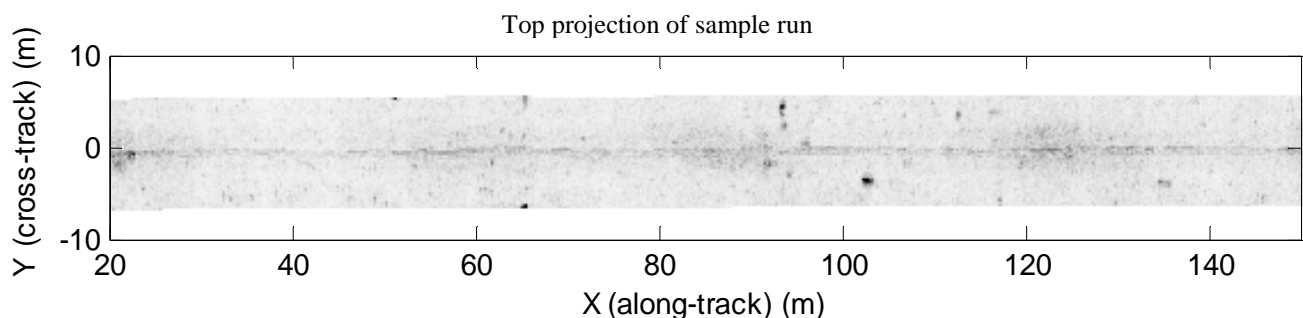
Top projection of sample run

Fig. 3. Top projection image of a sample BOSS SAS run. Note the cross-track extent of the data follows the path of the vehicle. The line down the center of the path is the specular return from the sea floor that has not been completely nulled.

This process is performed for a combination of vertical, forward, and aft look directions by applying overlapping, triangular spectral windows. Voxel energies from multiple look angles can be combined in a variety of manner to create a multi-aspect image – where in this work, the look angle yielding the highest energy value is used for each voxel.

Mosaicking and Geolocation

Processing data in elevation planes naturally results in image data represented in a cylindrical coordinate system. Linear interpolation between all the elevation beams and ranges is used to transform the data to Cartesian coordinates for subsequent display and analysis. The processor output at each along-track position is a vertical slice of backscatter image amplitudes, oriented perpendicular to the vehicle trajectory. Each of these 2-D vertical image slices are tagged with latitude/longitude, depth, and heading information.

The vertical slices can be mosaicked into a three-dimensional cube (see Fig. 3 for a sample output). This is achieved by traversing the data set to determine the physical extent of the data and to determine the best fit volume containing the data. The final volume of data contains the sonar data along with position and orientation information that allows the volume to be geolocated.

Multi-frame Processing

To enable continuous, streaming processing, the input pings are segmented into collections of sequential pings (a frame) to be independently SAS processed. The length of a frame is typically chosen to be between one and three synthetic apertures. Each frame uses an independent best-fit line during motion compensation. When a frame's processing is complete the vertical slices of SAS data are output. If frame overlap is used (to preserve the quality of data at the edges of the frame due to incomplete phase history in the azimuth direction), the vertical slices output by the processor are not necessarily in order spatially. However, each slice is tagged with the appropriate position and orientation information to construct the mosaicked image correctly. Additional frame overlap is necessary to avoid gaps in the final data when the vehicle makes significant course changes.

III. DATA AND RESULTS

The 3-D SAS processor and multi-aspect algorithm were successfully tested using a variety of data sets. Fig. 4 illustrates the beam steering and capabilities of the processor. A cylinder shaped object of approximate dimensions 1.5 m by 0.4 m is oriented slightly off horizontal (sloped upward along the direction of travel). The object is not highly visible when the synthetic receive beam is oriented straight down; however, steering the look direction normal to the object surface results in a squinted view that highlights the dominant specular returns. When the eleven different aspects are combined incoherently, the resulting multi-aspect image reveals the overall structure of the target.
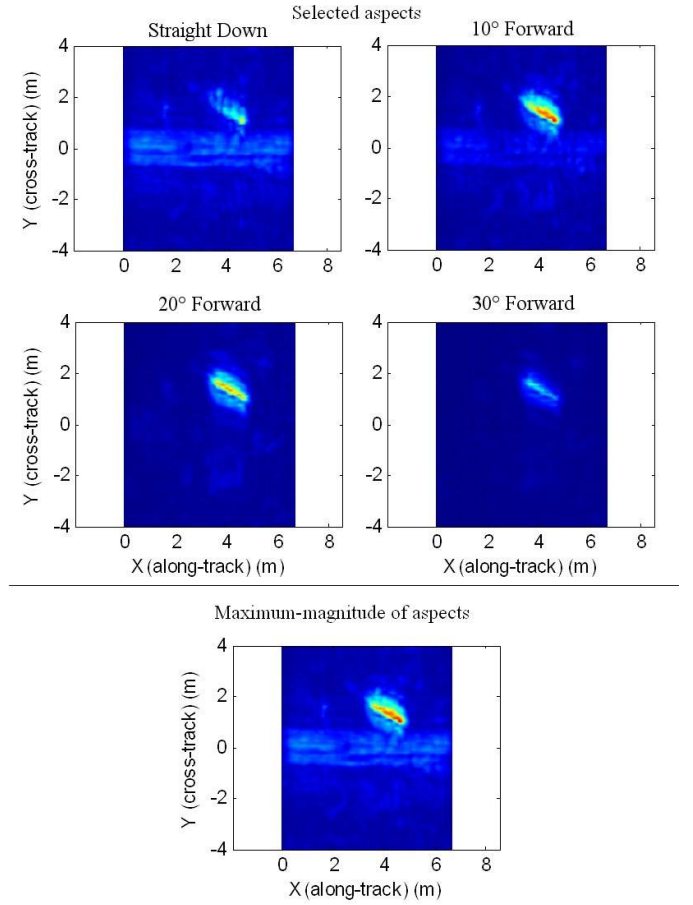

Fig. 4. Top projection images of a cylindrical object. Vehicle trajectory left to right along zero center line. Images above line are results of squinted SAS processing. Image below line is the multi-aspect result of combining 11 aspects together. The object slopes upward toward center line.
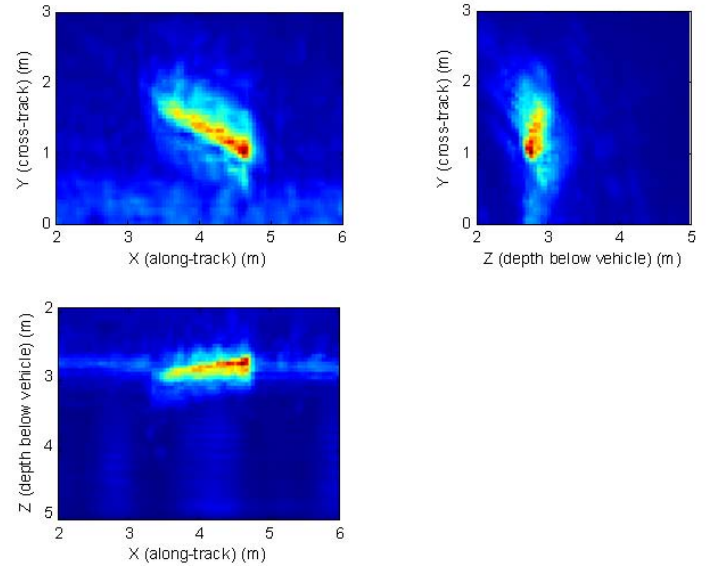

Fig. 5. Multi-aspect SAS Image of Cylinder: Top, front and side view projections.

Top, front, and side view projections of the multi-aspect image are presented in Fig. 5, verifying that fine along-track resolution is achieved. A receiver bandwidth of 8 kHz (~10 cm range resolution) was used to create this image. Magnitude plots through the strongest reflectors estimate an along-track point response consistent with 10 cm along-track resolution.

Efficiency of Frequency Domain 3-D SAS Beamformer

One of the benefits of utilizing the frequency domain beamformer for the along-track image focusing is the inherent algorithmic efficiency of this approach. The same speed benefits realized by frequency domain processing in traditional sidescan SAS apply to each elevation beam individually as it is essentially the same processing performed. A limit to this speedup may be encountered when a very small range extent is processed. In this case, the setup overhead of the frequency domain processor might become significant. However, as the problem size grows, the frequency domain approach described in this paper will grow in efficiency relative to the time-domain approach.

An additional benefit of the frequency domain approach pertains to the ease and algorithmic efficiency of generating multiple-aspect imagery. All motion compensation and elevation beamforming, and the majority of the RMA algorithm is computed once, and only the last operation (inverse FFT) is repeated multiple times with a different Doppler window to gain the independent looks of the scene. The amount of effort required to generate individual aspects is minimal as each aspect shares the majority of the computations with the other look directions.

IV. SUMMARY

The frequency domain approach to performing 3-D SAS processing described in this paper is computationally more efficient than the more common time-domain method.

The algorithm implemented uses traditional near-field, 2-D time-domain beamforming to generate elevation beams (planes) which extend in the along-track direction. Each of these elevation beams is independently SAS processed using the range migration algorithm. The resulting beamformed planes are then transformed onto a Cartesian grid for display and analysis.

The inherent algorithmic efficiency of RMA versus traditional 3-D time-domain beamforming results in improved overall data processing throughput. Generating aspect-dependent imagery requires little additional computational effort and is a natural byproduct of the frequency-domain approach.

Follow-on improvements to this processor include further optimizing the algorithmic efficiency. Because the elevation beamforming step still requires time-domain beamforming, a natural improvement would be a full 3-D frequency-domain SAS processor. This could further improve the efficiency gains realized by the hybrid time-domain/frequency-domain approach described here.

REFERENCES

[1] S. G. Schock, "Synthetic aperture 3D buried object imaging," Int. J Soc Underwater Tech., 27(4), pp. 185–193, 2008.
[2] M. Rognstad, "Subbottom swath mapping with HAWAII MR1," MTS/IEEE OCEANS 1996, pp. 1014-1018, Sep 1996.
[3] C. Cafforio, C. Prati, and F. Rocca, "SAR Data Focusing Using Seismic Migration Techniques," IEEE Transactions on Aerospace and Electronic Systems, vol. 27, pp. 194-206, 1991.